arXiv:1701.07434v1 [cs.NI] 25 Jan 2017

Asynchronous iterations in ultrametric spaces

Alexander J. T. Gurney

January 27, 2017

Abstract

Some iterative calculations can be carried out by parallel communicating processors, and yield the same results whether or not the processors are synchronized. We show that this is the case if and only if the iteration is a contraction that is strict on orbits, with respect to an ultrametric defined on the state space. The maximum number of independent processors is given by the dimension of the space.

We apply this theorem to interdomain routing, and are able to provide two advances over the previous state of the art. Firstly, multipath routing problems have unique solutions, if certain conditions are satisfied that are analogous to known correctness conditions for the single-path case. Secondly, these solutions can be computed asynchronously in a variety of ways, which go beyond methods that are currently used.

1 Introduction

The theory of *asynchronous iterations* is concerned with the problem of when an iterative algorithm can be implemented on a set of communicating processors, without explicit synchronization, and yet still compute the same result. It is known that for this to be possible, certain characteristics of the iteration must hold with respect to its state space: several different sufficient conditions are known. These are special cases of a more general result, which gives a necessary and sufficient condition for asynchronous safety. It requires that the state space have a 'nested box' structure, and that synchronous iterations always lead to a more-inward box. This condition is rather 'low-level', and may be difficult to verify in many cases; equally, the various sufficient conditions are easier to work with, but do not account for all possibilities. In addition, much of the prior work on these iterations assumes that data values are real numbers, whereas there are many iterative algorithms that work over other kinds of data.

In this paper, we reinterpret the 'nested box' structure in terms of a special kind of metric on the state space. In such an *ultrametric*, the balls around a given point always form nested boxes. The Banach fixed point theorem for ultrametric spaces, which states that a self-map of the space that is contractive and strict on orbits must have a unique fixed point, is precisely the theorem needed to prove asynchronous safety. That is, the application of this theorem proves not only that there is a unique fixed point, but also that it can be found by asynchronous iteration. Conversely, whenever an iteration is asynchronously safe, then an ultrametric can be defined with respect to which the iteration is a contraction of the required kind. Furthermore, the degree of potential asynchrony (the number of processors across which the iteration can be partitioned) is given by the dimension of the ultrametric space. This result applies to discrete data as well as to numeric problems.

In the final part of the paper, we apply this new theorem to a problem in interdomain multipath routing. The presence of a unique fixed point, let alone the possibility of asynchronous implementation, was not previously known. Existing sufficient conditions did not cover this case, but it is dealt with by the new theorem.

2 Background

This section explains the two separate areas of theory—asynchronous iterations (Section 2.1) and ultrametric spaces (Section 2.2)—which are involved in the main results of this paper.

2.1 Asynchronous iterations

There are many algorithms which operate by iteratively applying the same function to some state vector. If the state space is $M = M_1 \times M_2 \times \cdots \times M_k$, then a function σ from M to M can be decomposed as the product of k functions $\sigma_i : M \to M_i$, where

$$\sigma(m_1, m_2, \ldots, m_k) = (\sigma_1(m_1, m_2, \ldots, m_k), \sigma_2(m_1, m_2, \ldots, m_k), \ldots, \sigma_k(m_1, m_2, \ldots, m_k)).$$

On a sequential machine, each σ_i function must be evaluated in turn in order to produce the new state vector. But if multiple processors are available, then each function evaluation can be assigned to a different processor; once they have all finished, they can mutually communicate their results so that the next iteration step can begin. The total execution time for a single iteration is thus bounded by the time taken for the slowest processor.

It has long been known that for *some* algorithms, synchronous execution is not required for convergence. That is, the same answer as produced by the above process can also be generated in a far less restrictive execution model. In the asynchronous execution of the same algorithm, the same function is executed on each processor, but the execution is no longer in lock-step and the input data may come from a prior round of the iteration.

Asynchronous iteration has been applied to many problems: finding shortest paths [22], dynamic programming [3], finding fixed points of linear and non-linear operators [1], PageRank [13], and several others [8]. Use of the method is often motivated by the large quantity of data involved, or by the difficulty in ensuring synchronized execution. Sometimes, particularly for numeric problems, convergence time can be improved by dropping the synchrony requirement. While the theory was originally developed for numeric iterations [1, 2, 5], results have also been obtained for iterations on discrete data [23].

In the following, we take time to be discrete and linear; the set T contains all time values.

Definition 1. For a set P of processors, an *asynchronous execution schedule* consists of two functions α and β, where

- $\alpha : T \to 2^P$ yields the set of processors which activate at each time step, and
- $\beta : T \times P \times P \to T$ yields the delay between two given processors at each time step; so if $\beta(t,i,j) = t'$, then the data from j used at i at time t was generated at time t'.

Definition 2. A schedule (α, β) on P is *admissible* if

1. For all i in P, and t in T, there exists $t' > t$ such that i is in $\alpha(t')$.
2. For all t in T, and i and j in P, $\beta(t,i,j) > t$.
3. For all i and j in P, and t' in T, there exists a t_f in T such that if $t > t_f$, then $\beta(t,i,j) \neq t'$.

These admissibility conditions may be expressed more informally, as: every node activates infinitely often; information does not propagate backwards in time; and a past data value can only be used finitely often. Some weaker versions of these conditions have also been considered; for example, the final axiom may be replaced by an upper bound on the age of any data item used in a calculation [23].

If σ is an function from M to M, then we can define an asynchronous iteration corresponding to σ for any given schedule, and for a particular starting point in M. For each i, there will be a series of values in M_i generated by the iteration; call these $x_i(t)$ for $t \in T$. Let $x(t)$ be their product, so $x(t)$ is a vector in M.

Let m be a point in M, and let (α, β) be a schedule on a set of k processors. For each i, let $x_i(0) = m_i$. For $t > 0$ we define $x_i(t)$ by

$$x_i(t) = \begin{cases} x_i(t-1) & i \notin \alpha(t) \\ \sigma_i\left(x_1(\beta(t,i,1)), x_2(\beta(t,i,2)), \ldots, x_k(\beta(t,i,k))\right) & i \in \alpha(t). \end{cases}$$

So if processor i does not update at time t, its value does not change; when it does update, it carries out its usual σ_i operation, but may use values from further in the past than the immediately preceding step.

Note that if $\alpha(t) = \{1, 2, \ldots, k\}$ for all t, and $\beta(t,i,j) = t-1$ for all t, i and j, then this is just the synchronous iteration from above.

Definition 3. Let $M = \prod_{i \in I} M_i$, for some sets M_i. Let σ be a function from M to M that has a unique fixed point m^*. Then σ is an *asynchronously contracting operator* (ACO) if, for any admissible schedule (α, β) on $|I|$, and any starting point m in M, there is some time $T_{m,\alpha,\beta}$ such that for any $t > T_{m,\alpha,\beta}$, the state $x(t)$ is equal to m^*.

These ACOs may also be characterized in terms of a 'nested box' structure on the state space. Informally, if the synchronous iteration is such that it always takes points into a more-inward box, then it is asynchronously safe. The asynchronous iterations will eventually lead to the same fixed point, but may deviate from the synchronous course of execution.

Definition 4. A subset N of $M = \prod_{i \in I} M_i$ is a *box* if, for each i, there is a subset N_i of M_i such that $N = \prod_{i \in I} N_i$.

Theorem 1. *An operator σ on $M = \prod_{i \in I} M_i$ is an ACO if and only if there exist boxes $\{C_0, C_1, \ldots, C_k\}$ in M with the following properties:*

1. *$C_0 = \{m^*\}$ for some m^* in M.*
2. *$C_k = M$.*
3. *If $0 \leq r < s \leq k$, then $C_r \subset C_s$.*
4. *If m is in C_{r+1}, then $\sigma(m)$ is in C_r; and if m is in C_0 then $\sigma(m)$ is in C_0.*

Proof. See [23]. □

This powerful theorem means that we can determine asynchronous correctness, without having to reason about asynchronous processes. We merely have to verify that the synchronous iteration converges, and that certain conditions hold for the state space. Unfortunately, these conditions may be rather tricky to apply in practice—particularly if one wants to demonstrate that an operator is *not* an ACO. A further problem is that the low-level way in which the conditions are stated makes it difficult to understand the class of ACOs in general.

Several sufficient conditions are known which imply that the criteria of Theorem 1 are fulfilled. Numeric examples include weighted maximum norms over Banach spaces [6, 8], P-contractions [1], paracontractions [7, 17], and isotone mappings [16]. For discrete data, there are also results about isotone mappings [23]. All of these impose further requirements on the state space and iteration; the situation is especially painful for iterations on discrete data, since none of the usual real-number apparatus is available: we do not have continuity, norms, or even subtraction. The result of Section 3 is a necessary and sufficient condition for iterations on discrete data, so it covers iterations that are not necessarily isotone but even so manage to converge to a unique fixed point.

2.2 Ultrametric spaces

An ultrametric is a particular kind of 'distance' measurement that differs in several important respects from more familiar examples, but which also has useful applications.

A conventional metric space allows one to measure the distance between two points as a real number, with certain intuitive properties being fulfilled: a zero distance means the points are identical, distances are symmetric, and the triangle inequality is valid. In an ultrametric space, the triangle inequality is strengthened. If l, m, and n are three points in the space, then they

must form an isosceles triangle: two of the distances $d(l,m)$, $d(m,n)$ and $d(n,l)$ are the same. Furthermore, the remaining distance can be no longer than the others, so the triangle is 'long and narrow' as opposed to 'short and wide'. For this reason, the spaces are sometimes called *isosceles spaces*.

Definition 5. An *ultrametric space* (M,d,Γ) consists of a set M, a totally ordered set Γ with least element 0, and a function $d : M \times M \to \Gamma$ such that

1. $d(m,n) = 0$ if and only if $m = n$
2. $d(m,n) = d(n,m)$
3. $d(l,n) \leq \max(d(l,m), d(m,n))$

for all l, m and n in M.

A canonical example is when M is the set of all strings over some alphabet. For distinct x and y in M, let

$$m(x,y) = \min\left\{i \in \mathbb{N} \;\middle|\; x_i \neq y_i\right\}$$

so m yields the first index at which the two strings differ. Then

$$d(x,y) = \begin{cases} 0 & x = y \\ 2^{-m(x,y)} & x \neq y \end{cases}$$

is an ultrametric distance function.

Indeed, this example is very close to being universal: any ultrametric space M is isometric to a space where the elements are functions from $\mathbb{Q}_{\geq 0}$ to M, and the distance is given by

$$d(f,g) = \sup\left\{q \in \mathbb{Q}_{\geq 0} \;\middle|\; f(q) \neq g(q)\right\}.$$

See [14] for more details. Some other examples come from Boolean algebra, where the distance between two elements can be defined in terms of their symmetric difference [20].

If h is a function from M to $\Gamma \setminus \{0\}$, then a distance function can be defined by

$$d_h(m,n) = \begin{cases} 0 & m = n \\ \max(h(m), h(n)) & m \neq n. \end{cases}$$

This is clearly an ultrametric.

Definition 6. In an ultrametric space (M,d,Γ), the *ball* about a point m in M, of radius r in Γ, is the set

$$B(m;r) = \{n \in M \mid d(m,n) \leq r\}.$$

The balls of an ultrametric space have some surprising properties. Any point in a ball will serve as its center (see the lemma below). If we have two balls $B(n;r)$ and $B(m;s)$, then either one is a subset of the other, or they are disjoint. Consequently, the set of all balls, ordered by inclusion, has a tree structure [15].

Lemma 2. *If n is in $B(m;r)$ then $B(n;r) = B(m;r)$.*

Proof. If n is in $B(m;r)$ then $d(m,n) \leq r$. Then for any z in M,

$$d(m,z) \leq \max(d(m,n), d(n,z))$$

so if $d(n,z) \leq r$ then $d(m,z) \leq r$ as well; and the same argument applies if m and n are interchanged. Hence the two balls have the same content. □

We will need some notion of completeness of an ultrametric space, in order to guarantee the existence of fixed points. Otherwise, it could be that for certain iterations, the sequence of values converges to a point that is not in the space. The following definition is sufficient to ensure that the Banach fixed-point theorem actually yields a fixed point.

Definition 7. An ultrametric space is *spherically complete* if every chain of balls has nonempty intersection.

Simple examples of spherically complete spaces include any space that is finite, and any for which the image of d is a finite subset of Γ. In both of these cases, any chain of balls is guaranteed to be finite, and its intersection is then equal to the smallest ball in the chain.

The Banach theorem for ultrametric spaces can be made to work for several different kinds of contracting operator. The general idea is that, with respect to the ultrametric distance, each application of the operator brings points closer together. Eventually, the entire space is contracted into a single point, which is the desired unique fixed point.

Definition 8. Let σ be a function from M to M. If (M,d,Γ) is an ultrametric space, then σ is (with respect to d):

- a *contraction* if $d(\sigma(m),\sigma(n)) \leq d(m,n)$ for all m and n in M
- a *strict contraction* if $d(\sigma(m),\sigma(n)) < d(m,n)$ for all distinct m and n in M
- a *strict contraction on orbits* if $d(\sigma(m),\sigma^2(m)) < d(m,\sigma(m))$, or $m = \sigma(m)$, for all m in M.

Note that if σ is a strict contraction, then it is necessarily a contraction that is strict on orbits.

Theorem 3. *If σ is a function from M to M, and (M,d,Γ) is a spherically complete ultrametric space with respect to which σ is a contraction that is strictly contracting on orbits, then σ has a unique fixed point.*

Proof. See [18] and [19]. □

We now demonstrate that ultrametric spaces can be combined via a product operation. Furthermore, in the resulting space, every ball is a box. This provides the desired connection with the theory of asynchronous iterations: the balls about the fixed point will be precisely the boxes demanded by the asynchronous iteration theorem.

Definition 9. Given ultrametric spaces (M_i,d_i,Γ) for $1 \leq i \leq k$, define the *ultrametric product space* (M,d,Γ) by

$$M = \prod_{1 \leq i \leq k} M_i$$
$$d(m,n) = \max_{1 \leq i \leq k} d_i(m_i,n_i)$$

where m and n are vectors in M.

We will refer to k as the *dimension* of M. This usage is appropriate for the case when M is a real or complex vector space. When M is discrete, its topological dimension is zero; but here, we will carry on using the term 'dimension' for the number of components of the product.

Lemma 4. *In an ultrametric product space, every ball is a box. That is, if (M, d, Γ) is the product of (M_i, d_i, Γ) for $1 \leq i \leq k$, then for any m in M and r in Γ,*

$$B(m; r) = \prod_{1 \leq i \leq k} B_i$$

where each set B_i is a subset of M_i.

Proof. For each i, let B_i be $B(m_i; r)$ in M_i. For any element x of M, we have:

$$\begin{aligned}
& x \in B(m; r) \\
\iff & r \geq d(x, m) \\
\iff & r \geq \max_{1 \leq i \leq k} d_i(x_i, m_i) \\
\iff & \forall i : 1 \leq i \leq k \implies r \geq d_i(x_i, m_i) \\
\iff & \forall i : 1 \leq i \leq k \implies x_i \in B(m_i; r) \\
\iff & \forall i : 1 \leq i \leq k \implies x_i \in B_i \\
\iff & x \in \prod_{1 \leq i \leq k} B_i.
\end{aligned}$$

□

3 The main result

This section is dedicated to proving the theorem below, which provides a necessary and sufficient condition for an operator to be an ACO, in terms of an ultrametric structure on the state space.

Theorem 5. *Let M be a set, and $\sigma : M \to M$ a function from M to M. Then σ is an asynchronously contracting operator on M if and only if there exists an ultrametric d on M, with finite image, and with respect to which σ is a contraction that is strict on orbits.*

We will prove the two directions of this theorem separately. In order to establish this result, we will need to use the property that an operator on M is asynchronously contracting if and only if a series of nested boxes in M exist, with certain properties. The existence of an ultrametric will provide these boxes, and conversely, given a series of boxes, we can define a suitable ultrametric.

Lemma 6. *If (M, d, Γ) is a ultrametric space, with respect to which σ is a contraction that is strict on orbits, and Γ is finite, then a series of boxes exists that has the required properties.*

Proof. From the theory of ultrametric spaces, the contraction conditions on σ provide that it has a unique fixed point in M; call this m^*. For every possible radius r in Γ, there is a ball of radius r about m^*:

$$B(m^*; r) = \left\{ m \in M \;\middle|\; d(m^*, m) \leq r \right\}.$$

These balls will be the required boxes.

Firstly, $B(m^*; 0) = \{m^*\}$, since no other points are at distance zero from m^* itself.

Next, due to finiteness of Γ, there must be some minimal radius k such that $B(m^*; k) = M$. Let R be the set $\{r \mid 0 \leq r \leq k\}$. Clearly, if $r_1 < r_2$ for some r_1 and r_2 in R, then $B(m^*; r_1) \subseteq B(m^*; r_2)$. Because some of these balls may coincide, despite having different radii, we define a subset S of R such that

1. if r is in R then there is some s in S such that $s \leq r$ and $B(m^*; r) = B(m^*; s)$; and

2. for any s_1 and s_2 in S, $B(m^*; s_1) \neq B(m^*; s_2)$.

This S now yields the required sequence of boxes; it is well-defined since R is finite.

It remains to show that σ fulfils the required property. For any m in M, we have the relationship

$$d(m, \sigma(m)) = d(m, m^*)$$

and, if m is not equal to $\sigma(m)$,

$$d(m, \sigma(m)) > d(\sigma(m), \sigma^2(m)).$$

Hence

$$d(m, m^*) > d(\sigma(m), m^*)$$

unless m and $\sigma(m)$ are equal (in which case they are both equal to m^* itself). Consequently, application of σ always takes a point to a more-inward ball, unless that point is the fixed point already. □

The next step is to prove the converse: that if an operator is asynchronously contracting, then there is an ultrametric with respect to which the operator is a contraction that is strict on orbits.

Lemma 7. *Let M be a set endowed with an asynchronously contracting operator σ. Then there is an ultrametric d on M with respect to which σ is a contraction that is strictly contracting on orbits.*

Proof. Since σ is an asynchronously contracting operator, there exists a series of nested boxes with certain properties. We will define an ultrametric distance function that uses this box structure.

Let $C_0, C_1, \ldots, C_k$ be the box sequence, where C_0 is a singleton set, $C_k = M$, and $C_i \subset C_j$ whenever $0 \leq i < j \leq k$.

For any point m in M, we can find the index of the innermost hypercube that contains m:

$$C(m) = \min\{i \mid m \in C_i\}.$$

Thus $C(m^*) = 0$ if and only if m^* is the unique fixed point. The required distance function is

$$d_C(m, n) = \begin{cases} 0 & m = n \\ \max(C(m), C(n)) & m \neq n \end{cases}$$

which is an ultrametric. Note that there are only finitely many possible radii.

We can now show that the balls about m^* with respect to d are precisely the given boxes C_i. For any radius r,

$$m \in B(m^*; r) \iff d(m, m^*) \leq r \iff C(m) \leq r \iff \min\{i \mid m \in C_i\} \leq r \iff m \in C_r.$$

Finally, we prove that σ must be a contraction that is strictly contracting on orbits. Note that $C(\sigma(m)) < C(m)$ for all m other than m^*. Hence $d(\sigma(m), m^*) < d(m, m^*)$ for such m. Since $d(m, m^*) = d(m, \sigma(m))$ for all m, we have

$$d(m, \sigma(m)) > d(\sigma(m), \sigma^2(m))$$

whenever $m \neq \sigma(m)$, so σ is strictly contracting on orbits. Similarly, if m and n are two points in M, then $C(\sigma(m)) \leq C(m)$ and $C(\sigma(n)) \leq C(n)$. Therefore the larger of $C(\sigma(m))$ and $C(\sigma(n))$ cannot exceed the larger of $C(m)$ and $C(n)$; we obtain

$$d(\sigma(m), \sigma(n)) \leq d(m, n)$$

as desired. □

This completes the proof. We have established that an operator is asynchronously contracting if and only if it fulfils the ultrametric Banach fixed point theorem. Furthermore, the degree of asynchrony is given by the dimension of the ultrametric space. This provides a convenient proof technique for asynchronous iterations, subsuming several other previously-known special conditions.

4 Recovery of previous theorems

The general theorem of Section 3 has many specific consequences for particular classes of iteration. Several of these have previously been studied in the literature. In this section, we relate previous results to the new theory, thereby demonstrating its generality.

4.1 Paracontractions

In the case of iterations over real vectors, there is a well-known theory of paracontracting operators. In the following, let M be $\mathbb{R}^n_{\geq 0}$, where $n \geq 1$. For v in M, let $\|v\|$ be the vector whose ith entry is the absolute value of v_i. For two vectors v and w in M, say that $v \leq w$ if for all i, $v_i \leq w_i$.

Definition 10. A function σ from M to M is called a *paracontraction* if there exists an n by n matrix P with entries in $\mathbb{R}_{\geq 0}$, having spectral radius less than 1, and for which

$$\left\|\sigma(x) - \sigma(y)\right\| \leq P \left\|x - y\right\|.$$

Theorem 8. *If σ is a paracontraction on M, then (M, d) is an ultrametric space with respect to which σ is a strict contraction, where*

$$d(x, y) = \max_{1 \leq i \leq n} \alpha_i \left\|x - y\right\|_i.$$

Proof. If σ is paracontracting, then it is a strict contraction with respect to a weighted maximum norm [1, 2]. So the previous result suffices to prove this one. □

4.2 Weighted maximum norms

Suppose that the set M is a product of sets M_1 through M_n, each of which is equipped with a real-valued norm $\|\cdot\|_i$. For any real numbers α_i, a norm can be defined on M by

$$\|x\| = \max_i \alpha_i \|x_i\|_i.$$

An operator σ on M is *Lipschitz* if there exists some p, with $0 \leq p < 1$, such that

$$\left\|\sigma(x) - \sigma(y)\right\| \leq p \left\|x - y\right\|$$

for all x and y in M.

Theorem 9. *If σ is Lipschitz with respect to a weighted max-norm, then it is also a strict contraction with respect to an ultrametric distance on M, of dimension n.*

Proof. For each i, let d_i be the ultrametric on M_i given by

$$d_i(x, y) = \alpha_i \left\| x - y \right\|_i .$$

□

4.3 Monotonic contractions

Let $\leq$ denote the direct product order on vectors in $(\mathbb{R}_{\geq 0})^n$, so that

$$x \leq y \iff \forall i : x_i \leq_i y_i.$$

A function σ on M is monotonic if

$$x \leq y \implies \sigma(x) \leq \sigma(y)$$

for all x and y in M. If σ is a monotonic function, with a unique fixed point x^*, and (other conditions) then it is an asynchronously contracting operator.

This conclusion also follows from the theorem of this paper. To construct the required ultrametric, we can define ultrametrics on each component, and then take their max-product. The individual ultrametrics can be defined in terms of the natural ordering on $\mathbb{R}_{\geq 0}$.

Let $h(x) = 2^{-x}$. Note that for x greater than or equal to zero, $h(x)$ is in the range $(0,\ 1]$. Therefore, the function

$$d(x, y) = \begin{cases} 0 & x = y \\ \max(h(x), h(y)) & x \neq y \end{cases}$$

is an ultrametric distance function on $\mathbb{R}_{\geq 0}$; and so the product of n of these is also an ultrametric, on $(\mathbb{R}_{\geq 0})^n$. The conditions on σ imply that it is contractive with respect to this ultrametric.

5 Application to interdomain multipath routing

We will now see an extended example of the use of the ultrametric theorem to prove asynchronous safety of an iteration. The iteration in question is simple to describe, but has some unusual properties which make it unsuitable for handling by previously-known asynchrony theorems.

The problem comes from the selection of paths at the interdomain level in network routing. The various networks which combine to make the Internet carry out path selection in a way which provides a great deal of local autonomy: the paths for a given source-destination pair could in principle be ranked arbitrarily. Because of the local nature of preferences and decisions, the overall routing outcome is not a global optimum, as for shortest-path algorithms, but is a Nash equilibrium between the networks involved [10]. This problem is also inherently distributed and asynchronous: the private nature of policy means that the computation cannot be performed centrally, and synchronization on a global scale is infeasible.

In execution of the path selection process, even synchronously, various anomalies appear which would be impossible for shortest path algorithms. It is *not* necessarily the case that the paths selected by a given node improve over time: a node is perfectly capable of switching to a

worse path if its previous path becomes unavailable. A path could be lost and then regained, possibly several times. Across the entire network, it may be that at a given time step *all* nodes either stay the same or are forced to choose a worse path. Nevertheless, an eventual fixed point can still (sometimes) be found, even after all of these strange events have occurred.

Theoretical models of this situation include the *stable paths problem* [10]. This is a combinatorial game where the Nash equilibria correspond to solutions of the interdomain routing problem. An instance of the stable paths problem is given by:

- A graph $G = (V, E)$ with a designated destination node d in V.
- For each v in V, a partial ranking of the simple paths in G from v to d. (That is, a total order on a subset of these paths.) These are called the permitted paths.

The empty path (from d to d) is always permitted. A solution of the problem is a stable path assignment. A *path assignment* associates each node in V with at most one of its permitted paths. A path assignment π is *stable* if, for each v other than d itself, $\pi(v)$ is identical to the most-preferred path in $\{(v\ w)\pi(w) \mid (v\ w) \in E\}$, and $\pi(d)$ is the empty path.

A stable path assignment is therefore a fixed point of the myopic best-response iteration. Stable path instances are known to exist with zero, one, or several stable solutions. Certain sufficient criteria for the existence of a unique stable solution are known, but may be NP-hard to check. The iteration has previously been shown to be asynchronously safe, with each node being responsible for managing its local state [10]. The proof relies on a message-passing model with explicit queues, rather than using the 'asynchronous iteration' framework. An alternative proof based on asynchronous iterations did not succeed in demonstrating the desired result [4].

In this section, we extend the stable paths problem to the selection of multiple paths. (There is a version of the stable paths problem in which mixed Nash equilibria are allowed: this is, however, problem from the multiple-path problem presented here [12].) The proof uses the theory developed in this paper. Consequently, this establishes a new correctness result for multiple stable paths; and the result easily extends to the asynchronous case. This example had previously been treated using ultrametric spaces, but without any proof of asynchronous correctness [11].

A known sufficient condition for existence of a unique solution, in the case of single-path selection, is that the path preferences be determined according to a strictly inflationary order [9, 21]. That is, there is a total order on simple paths, such that a path is strictly preferred to any extension of that same path.

For multiple paths, we can make a similar definition. Suppose that the set $\mathscr{P}$ of simple paths to d has a preorder $\preceq$, with $p \prec q$ indicating that p is preferred to q. (Recall that a preorder is a reflexive and transitive relation.) If, for any path p from j to k, and any arc $(i\ j)$, we have $p \prec (i\ j)p$, then the preferences are *strictly inflationary*. For any subset A of $\mathscr{P}$, we can define

$$\min(A) = \{a \in A \mid \forall b \in A : \neg(b \prec a)\},$$

the set of minimal elements in A.

We can now define the iteration. The state space M is a product of sets M_i for i in V. Each M_i is the powerset of the set of simple paths from i to d. Therefore, M is isomorphic to the powerset of $\mathscr{P}$. The iteration proceeds as follows:

$$\sigma(x)_i = \begin{cases} \min\left\{(i\ j)p \;\middle|\; (i\ j) \in E, p \in x_j\right\} & i \neq d \\ \{\epsilon\} & i = d \end{cases}$$

where ϵ is the empty path. So at each step, a node sees the paths which have been selected by its neighbors, and from their extensions, chooses its 'best' paths. Note that a node may *lose* a path if its prefix ceases to be chosen by a neighbor.

The distance function for the space is defined by

$$d(m,n) = \max\left\{h(p) \;\middle|\; p \in (m \setminus n) \cup (n \setminus m)\right\}$$

for elements m and n of M, where $h(p) = \left|\left\{q \in \mathscr{P} \;\middle|\; q \geq p\right\}\right|$. Note that if $p \prec q$, then $h(p) > h(q)$ If the set $m \setminus n \cup n \setminus m$ is empty, then consistently with the definition of max, we take $d(m,n) = 0$. This gives an ultrametric space.

To prove that the best-path selection process always terminates in this case, and that it is asynchronously safe, we need only show that σ is a strict contraction.

Theorem 10. *With M, d and σ as defined above, σ is a strict contraction.*

Proof. We must show that

$$d(\sigma(m),\sigma(n)) < d(m,n)$$

for all distinct m and n in M. If $\sigma(m) = \sigma(n)$ then there is nothing to prove, so assume that they are different. Then

$$d(\sigma(m),\sigma(n)) = h(p)$$

for some p; without loss of generality, p is in $\sigma(m)$ but not $\sigma(n)$.

Note that p cannot be the empty path, since this will be present in both $\sigma(m)_d$ and $\sigma(n)_d$. Therefore $p = (i\ j)q$ for some q in m_j. By the strict inflationary property, we have $q \prec p$, and hence $h(q) > h(p)$.

Now, we will show that q is not in n_j. If it was, then p would have been a candidate path for $\sigma(n)_i$; but it was not selected. The only reason for that to happen would be that $\sigma(n)_i$ chose some better path p' instead. That p' could not be in $\sigma(m)_i$, since $\sigma(m)_i$ does contain a worse path, namely p; so p' is in the symmetric difference of $\sigma(m)$ and $\sigma(n)$. But if $p' \prec p$, then $h(p') > h(p)$, contradicting the choice of p as having greatest h value among all paths in the symmetric difference.

Therefore, q is not in n_j. Since it is in the symmetric difference of m and n, we have

$$d(\sigma(m),\sigma(n)) = h(p) < h(q) \leq d(m,n)$$

which establishes σ as a strict contraction. □

This result demonstrates, quite succinctly, that the iterative path-finding process always finds its unique fixed point, and that this iteration can be implemented asynchronously.

It is immediate from the definition that the ultrametric space (M,d) can be written as a product in several ways. Each of these corresponds to a mode of implementing σ by separate, asynchronously communicating processors. Indeed, the maximal possible decomposition for which σ remains an ACO is for the presence or absence of *each individual path* to be calculated by a separate processor:

$$(M,d) = \prod_{p \in \mathscr{P}} (M_p, d_p)$$

where $M_p = \{\epsilon, p\}$ and $d_p(\epsilon, p) = h(p)$.

This is significantly more relaxed than current execution models for interdomain routing, which admit asynchrony to the extent that a node may compute paths to all of its destinations, separately from any other node computing paths to all of its destinations. The new proof demonstrates that we could relax this model in several ways. The unit of computation could be:

- Paths for a given source (as now).
- Paths for a given source and destination.
- Paths for a given source, destination, and next-hop (at the router level or at the autonomous system level).
- Paths originating from a given geographic region, or a given class of autonomous systems.
- Paths of a given level of preference.

In many of these cases, we admit the possibility that intermediate states of the routing system could contain 'inconsistencies'. For example, it could be that paths p and p', for the same source and destination, with p preferred to p', might be simultaneously present in the system state. While this may seem unusual, we have proved that it will not harm eventual convergence to a consistent state. Already, in existing routing, we consider it normal for a node to be using a path that is transiently inconsistent with its next-hop neighbor's choice: this is essentially the same idea. (Note that while the proof of Theorem 10 relies on such states not occuring, this result applies to the synchronous execution only: asynchronous executions are allowed to behave more wildly.)

We have therefore revealed that routing protocol designs can be extended in two ways, compared to the present state of the art. Firstly, as long as the above rules are followed, we can move from single-path to multipath routing with confidence. Secondly, the computation of routes by asynchronous processes can be accomplished in many other ways from the method previously known to be safe.

6 Application to logic programming

Consider logic programs consisting of clauses

$$A \leftarrow L_1, \ldots, L_n$$

where each A is an atom, and each L_i is either an atom or the negation of an atom. (There may be no atoms in the list, in which case the clause is called a 'fact'.)

Denote the Herbrand base of a program P by B_P.

A program P is said to be *locally stratified* if there is a mapping $\rho : B_P \longrightarrow \Gamma$ into a totally ordered set Γ, such that for each clause $A \leftarrow L_1, \ldots L_n$ in P, we have

1. $\rho(A) \geq \rho(B)$ for all positive atoms B, and
2. $\rho(A) > \rho(B)$ for all negative atoms B.

An *interpretation* of P is a map I from B_P to {true, false}.

The *immediate consequence operator* T_P is defined as (TODO).

Define an ultrametric d on interpretations by

$$d(I,J) = \min\left\{\rho(A) \;\middle|\; A \in (I \setminus J) \cup (J \setminus I)\right\}.$$

This is spherically complete, and T_P is a strict contraction on it. This demonstrates the existence of a unique fixed point, corresponding to the perfect model semantics of the program.

By recourse to the above theorem, we have also shown that execution of the logic program can be spread among independent processors. Indeed, any term in the Herbrand base could be the responsibility of a separate processor.

7 Conclusion

This paper has established a connection between the theories of asynchronous iterations and ultrametric spaces. It provides a complete characterization of asynchronously contracting operators, and hence gives new proof techniques by which asynchronous safety can be proved (or disproved) in particular cases.

It is hoped that this formalism will be useful in understanding the class of ACOs, and in finding new examples, particularly those which are not numeric. There is also work to be done in treating more complex families of iteration, such as those possessing multiple fixed points. The study of ultrametric embeddings and representations may also provide insight into asynchronous contractions and their properties.

In the realm of networking practice, the suggestion of alternative routing architectures is tantalising. Further research will be needed in order to establish which such designs are desirable.

References

[1] Gérard M. Baudet. Asynchronous iterative methods for multiprocessors. *J. ACM*, 25(2):226–244, 1978.

[2] Dimitri P. Bertsekas. Distributed asynchronous computation of fixed points. *Math. Programming*, 27:107–120, 1983.

[3] Dimitri P. Bertsekas and John N. Tsitsiklis. *Parallel and distributed computation: Numerical methods*. Prentice-Hall, 1989.

[4] Chi-Kin Chau. Policy-based routing with non-strict preferences. In *Proc. ACM SIGCOMM*, pages 387–398, 2006.

[5] Daniel Chazan and Willard L. Miranker. Chaotic relaxation. *Linear Algebra Appl.*, 2(2):199–222, 1969.

[6] Mouhamed Nabih El Tarazi. Some convergence results for asynchronous algorithms. *Numer. Math.*, 39(3):325–340, 1982.

[7] Ludwig Elsner, Israel Koltracht, and Michael Neumann. Convergence of sequential and asynchronous nonlinear paracontractions. *Numer. Math.*, 62(1):305–319, 1992.

[8] Andreas Frommer and Daniel B. Szyld. On asynchronous iterations. *J. Comput. Appl. Math.*, 123(1–2):201–216, 2000.

[9] Timothy G. Griffin and Alexander J. T. Gurney. Increasing bisemigroups and algebraic routing. In *Relations and Kleene algebra in computer science*, number 4988 in Lecture Notes in Computer Science, pages 123–137. Springer-Verlag, 2008.

[10] Timothy G. Griffin, F. Bruce Shepherd, and Gordon Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, 10(2):232–243, 2002.

[11] Alexander J. T. Gurney. *Construction and verification of routing algebras*. PhD thesis, University of Cambridge, 2009.

[12] Penny E. Haxell and Gordon Wilfong. A fractional model of the border gateway protocol (BGP). In *Proc. ACM-SIAM SODA*, pages 193–199, 2008.

[13] Giorgios Kollias, Efstratios Gallopoulos, and Daniel B. Szyld. Asynchronous iterative computations with web information retrieval structures: the PageRank case. In *Proc. ParCo*, pages 309–316, 2005.

[14] Alex J. Lemin. On ultrametrization of general metric spaces. *Proc. Amer. Math. Soc.*, 131(3):979–989, 2002.

[15] Alex J. Lemin. The category of ultrametric spaces is isomorphic to the category of complete, atomic, tree-like and real graduated lattices LAT*. *Algebra Universalis*, 50(1):35–49, 2003.

[16] Jean-Claude Miellou. Itérations chaotiques à retards; études de la convergence dans le cas d'espaces partiellment ordonnés. *C. R. Acad. Sci. Paris Sér. A*, 280:233–236, 1975.

[17] Matthias Pott. On the convergence of asynchronous iteration methods for nonlinear paracontractions and consistent linear systems. *Linear Algebra Appl.*, 283(1):1–33, 1998.

[18] Sibylla Prieß-Crampe. Der Banachsche Fixpunktsatz für ultrametrische Räume. *Results Math.*, 18(1–2):178–186, 1990.

[19] Sibylla Prieß-Crampe and Paolo Ribenboim. Results on fixed points, combs and generalized formal power series. *C. R. Acad. Sci. Paris Sér. I Math.*, 315(6):639–643, 1992.

[20] Sibylla Prieß-Crampe and Paolo Ribenboim. Generalized ultrametric spaces I. *Abh. Math. Sem. Univ. Hamburg*, 66:55–73, 1996.

[21] João Luís Sobrinho. An algebraic theory of dynamic network routing. *IEEE/ACM Trans. Netw.*, 13(5):1160–1173, 2005.

[22] Paul Tseng, Dimitri P. Bertsekas, and John N. Tsitsiklis. Partially asynchronous parallel algorithms for network flow and other problems. *SIAM J. Control Optim.*, 28(3):678–710, 1990.

[23] Aydin Üresin and Michel Dubois. Parallel asynchronous algorithms for discrete data. *J. ACM*, 37(3):588–606, 1990.